

AD-A158 324

DENSE TRIVALENT GRAPHS FOR PROCESSOR INTERCONNECTION(U)

1/1

WISCONSIN UNIV-MADISON W LELAND ET AL. 1981

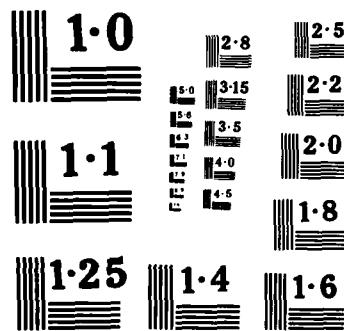
N00014-81-C-2151

UNCLASSIFIED

F/G 12/1

NL

END
FILED
STIC



NATIONAL BUREAU OF STANDARDS
MICROCOPY RESOLUTION TEST CHART

(1)

DENSE TRIVALENT GRAPHS FOR PROCESSOR INTERCONNECTION*

by

Will Leland and Marvin Solomon
University of Wisconsin--Madison

1981

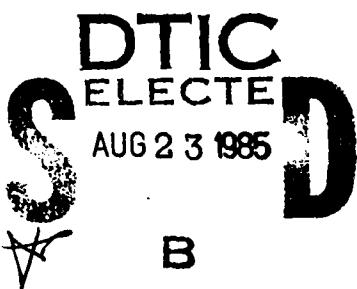
AD-A158 324

ABSTRACT

This paper presents a new family of undirected graphs that allows N processors to be connected in a network of diameter $3/2 \log_2 N + O(1)$, while only requiring that each processor be connected to three neighbors. The best trivalent graphs previously proposed require a diameter of $2 \log_2 N + O(1)$.

Key words and phrases: interconnection topology, multicomputer, undirected graphs, cubic graphs

DTIC FILE COPY



*This work was supported in part by the Defense Advanced Projects Research Agency under Contract No. N00014-81-C-2151.

1. INTRODUCTION

In the design of a network of processors, one important consideration is the interconnection topology. A multicomputer is a collection of processing elements (processor-memory complexes) connected by a communications medium. The communications medium may be shared memory (such as [1]), a broadcast bus [2] or ring [3], or a collection of point-to-point communications lines [4]. In this paper, we will be concerned with point-to-point networks.

Such a network can be modeled by an undirected simple graph [5] in which the vertices represent processing elements and the edges represent (full duplex) communication lines. To be useful as an interconnection graph, a graph should have the following properties:

1. The degree (maximum number of edges incident on any vertex) should be small. This restriction represents the fact that a processing element with a large number of line interfaces is expensive. If different vertices have different degrees, either all processing elements will need to be able to accommodate the maximum number of lines, or more than one kind of processing element will need to be designed and constructed.
2. The graph should be dense. In other words, the diameter (maximum distance between any pair of vertices, where the distance is the number of edges in a minimal path connecting the vertices) should be small relative to the number of

vertices. Processing elements not directly connected by a line will need to have messages relayed by intervening vertices. The number of relays in the worst case should be kept small.

- No vertex should be on an unusually high proportion of the shortest paths. To avoid congestion, the message-relaying load should be uniformly distributed throughout the network. If many pairs of vertices have shortest paths involving one central vertex, then that vertex must accommodate an especially large amount of traffic. As in point 1, we would be faced with a choice between designing all vertices to handle the worst case, or producing more than one kind of processing element.

The extreme case of property 1 is to allow each processor to have at most two neighbors. In this case, the best possible topology is a ring. Each vertex of a ring lies on the same number of shortest paths, but the diameter of a ring is equal to half the number of vertices.

Several authors have investigated families of trivalent (sometimes called cubic) graphs, in which each vertex has at most three neighbors. Arden and Lee [6] have introduced a family of graphs called chordal ring networks. The N vertices are arranged in a ring using two edges apiece. The third edge from a vertex is a chord connecting it to a vertex \sqrt{N} of the way around the ring. Arden and Lee show that the diameter is $O(\sqrt{N})$.

The cube-connected cycles of Preparata and Vuillemin [7] are trivalent graphs with $N = n \cdot 2^n$ vertices (for any integer n). The

<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Codes
1/or

Dist	special
A-1	



vertices are grouped into 2^n cycles of n vertices each, conceptually arranged at the corners of an n -dimensional cube. Each of the n edges emanating from a corner of the cube is used to connect one of the vertices of the corresponding cycle to a vertex of a neighboring cycle. Another way of describing the cube-connected cycles is to label each vertex with a pair (x, i) , where $x = x_0 \dots x_{n-1}$ is a sequence of n 0's and 1's and i is an integer in the interval $[0, n-1]$. The neighbors of vertex (x, i) are $(x, (i+1) \bmod n)$, $(x, (i-1) \bmod n)$, and (y, i) , where $y = x_0 \dots x_{i-1} \bar{x}_i x_{i+1} \dots x_{n-1}$. The diameter of the cube-connected cycles is $5/2 \lg N + O(1)$ (where \lg denotes the base 2 logarithm).

Another family of trivalent graphs is based on binary trees. More precisely, take three balanced binary trees of depth n , add a new vertex, and connect it to the roots of the trees. The resulting graph has $N = 3 \cdot 2^{n+1} - 2$ vertices and diameter $2n + 2 = 2 \lg N + O(1)$. Although this graph is denser than the chordal ring and the cube-connected cycles, it has severe congestion, with over half of all shortest paths traversing the root vertex.

A generalization of the binary tree, called multi-tree structured graphs [8], is formed by taking t binary trees and connecting the root of each to one vertex of a t -vertex cycle. The leaves of all the trees are then linked together in another cycle. Different graphs are produced depending on the choice of the sizes of the trees, the number of trees, and the order in which the leaves appear on their cycle. Arden and Lee show that if $t < \sqrt{N}$, a fairly simple interleaving of vertices can be used to guarantee a diameter of $2 \lg N + O(1)$. They use a heuristic

search program to find better interleavings for particular values of N and t , and conjecture, extrapolating from these results, that a diameter of $\lg N$ is possible for arbitrarily large N . However, they present no general algorithm for achieving this diameter.

In this paper, we introduce a new family of trivalent graphs with $N = 2^n$ vertices (for any $n \geq 2$) and diameter bounded by $3/2 \lg N$. These graphs also have much less congestion than binary trees or multi-tree structures.

The remainder of this paper is organized as follows: Section 2 introduces notation and defines the new family of graphs, which we call Moebius graphs. Section 3 proves the bound on the diameter. Section 4 summarizes our results and states some open problems.

2. DEFINITIONS

Let n denote a fixed integer greater than 1. Let 2^n denote the set $\{s = s_0 \dots s_{n-1} \mid s_i \in \{0,1\}\}$.

Conventions Throughout the remainder of this paper, all subscripts are to be interpreted modulo n . The symbol \oplus denotes addition modulo 2. Throughout section 3, Σ denotes summation modulo 2; when not specified otherwise, the sum is over all i , $0 \leq i < n$. The weight of s , denoted $w(s)$, is the number of 1-bits in s .

Define two functions f and g (mapping $2^n \rightarrow 2^n$) as follows:

$$f(s_0 \dots s_{n-1}) = s_1 \dots s_{n-1} \bar{s}_0$$

$$g(s_0 \dots s_{n-3} s_{n-2} s_{n-1}) = s_0 \dots s_{n-3} \bar{s}_{n-2} \bar{s}_{n-1}$$

Clearly, f and g are permutations of 2^n and $g = g^{-1}$. (The function f is often called the shuffle exchange [9].) Let Id denote the identity function on 2^n .

The Moebius graph of order n (so named because the function f introduces a loop with a "twist") is the graph $G = (V, E)$ with vertex-set $V = 2^n$ and $(u, v) \in E$ iff $u = f(v)$ or $v = f(u)$ or $u = g(v)$ (equivalently, $v = g(u)$).

Any path of vertices v^0, \dots, v^k in G can be described by a sequence $p = p_1 \dots p_k$, where for $i = 1, \dots, k$, $p_i \in \{f, f^{-1}, g\}$ and $v^i = p_i(v^{i-1})$. We will use the term "path" to mean the sequence of vertices, the sequence of edges connecting them, or the corresponding sequence of functions. No confusion should arise. We will write $v^k = p(v^0)$.

3. DIAMETER OF THE MOEBIUS GRAPH

In this section we will show that the diameter of the Moebius graph is bounded by $\lfloor 3/2 n \rfloor$. In fact, we obtain this result by proving the slightly stronger result that any pair of vertices is connected by a path of length $\leq \lfloor 3/2 n \rfloor$ that only traverses the f edges in the "positive" direction. More precisely, for any two vertices s and d , there is a path from s to d of the form $p_1 \dots p_k$, where $k \leq \lfloor 3/2 n \rfloor$ and $\forall i, p_i \in \{f, g\}$. Since $g^2 = Id$, we can, without loss of generality, restrict our attention to paths of the form $g_0 f g_1 f \dots f g_k$, where $\forall i, g_i \in \{g, Id\}$. We introduce

the following notation for such paths: Let $x = x_0 \dots x_{k-1} \in 2^k$. Then \hat{x} is the path $fg_0fg_1\dots fg_{k-1}$ and \tilde{x} is the path $g_0fg_1\dots g_{k-1}$, where

$$g_i = \begin{cases} g & \text{if } x_i = 1 \\ \text{Id} & \text{if } x_i = \emptyset \end{cases}$$

Lemma 1 If $a \in 2^1$ and $s = s_0 \dots s_{n-1}$, then $\hat{a}(s) = t = t_0 \dots t_{n-1}$,

$$\text{where } t_i = \begin{cases} s_{i+1} & \text{for } i \leq n-3 \\ s_{n-1} \oplus a & \text{for } i = n-2 \\ s_0 \oplus a \oplus 1 & \text{for } i = n-1 \end{cases}$$

Proof Clear from the definitions.

Lemma 2 If $x = x_0 \dots x_{n-1}$, $s = s_0 \dots s_{n-1}$, $d = d_0 \dots d_{n-1}$, and $d = \hat{x}(s)$, then for $i = 0, \dots, n-1$, $d_i = s_i \oplus x_i \oplus x_{i+1} \oplus 1$.

Proof For $k = 0, \dots, n$, let $x^k = x_0x_1\dots x_{k-1}$ and let $s^k = \hat{x}^k(s)$ (so x^0 is the empty path, $s^0 = s$, and $s^n = d$).

A tedious but straightforward proof by induction shows that for $1 \leq k \leq n-1$,

$$s_i^k = \begin{cases} s_{k+i} & \text{for } 0 \leq i \leq n-k-2 \\ s_{n-1} \oplus x_0 & \text{for } i = n-k-1 \\ s_{k+i} \oplus x_{k+i} \oplus x_{k+i+1} \oplus 1 & \text{for } n-k \leq i \leq n-2 \\ s_{k-1} \oplus x_{k-1} \oplus 1 & \text{for } i = n-1 \end{cases} \quad (1)$$

Now consider $d = s^n = \hat{x}_{n-1}(s^{n-1})$:

If $i \leq n-3$, then

$$\begin{aligned}
 s_i^n &= s_{i+1}^{n-1} \text{ by Lemma 1} \\
 &= s_{n-1+i+1} \oplus x_{n-1+i+1} \oplus x_{n-1+i+2} \oplus 1 \text{ by (1),} \\
 &\quad \text{since } n-(n-1) = 1 \leq i+1 \leq n-2 \\
 &= s_i \oplus x_i \oplus x_{i+1} \oplus 1
 \end{aligned}$$

If $i = n-2$, then

$$\begin{aligned}
 s_i^n &= s_{n-i}^{n-1} \oplus x_{n-1} \text{ by Lemma 1} \\
 &= s_{n-2} \oplus x_{n-2} \oplus 1 \oplus x_{n-1} \text{ by (1)} \\
 &= s_i \oplus x_i \oplus x_{i+1} \oplus 1
 \end{aligned}$$

If $i = n-1$, then

$$\begin{aligned}
 s_i^n &= s_{\emptyset}^{n-1} \oplus x_{n-1} \oplus 1 \text{ by Lemma 1} \\
 &= s_{n-1} \oplus x_{\emptyset} \oplus x_{n-1} \oplus 1 \text{ by (1)} \\
 &= s_i \oplus x_i \oplus x_{i+1} \oplus 1
 \end{aligned}$$

This completes the proof of Lemma 2.

Lemma 2 is illustrated for the case $n = 4$ in Figure 1.

Corollary 1 If $\sum(d_i \oplus s_i \oplus 1) = \emptyset$ then there are two strings $x, y \in 2^n$ such that $\hat{x}(s) = \hat{y}(s) = d$ and $x = \bar{y}$ (i.e., $x_i = y_i \oplus 1$ for $i = \emptyset, \dots, n-1$).

Proof For fixed s and d , Lemma 2 provides n linear equations in the n unknowns x_i :

$$x_{i+1} = d_i \oplus s_i \oplus 1 \oplus x_i, \quad \emptyset \leq i < n \quad (2)$$

However, these equations are not independent. Solving in terms of x_{\emptyset} yields

$$x_i = x_{\emptyset} \oplus \sum_{j < i} (d_j \oplus s_j \oplus 1), \quad 1 \leq i < n$$

$$\text{and } x_{\emptyset} = x_{\emptyset} \oplus \sum_{i < n} (d_i \oplus s_i \oplus 1),$$

so (2) has a solution iff $\zeta(d_i \oplus s_i \oplus 1) = 0$. If this condition is satisfied, the assignments $x_0 = 0$ and $x_0 = 1$ yield two solutions which are complements of each other.

Corollary 2 If $\zeta(d_i \oplus s_i \oplus 1) = 0$, then there exists a path from s to d of length less than or equal to $|3/2 n|$.

Proof The length of any path \hat{x} is $n + w(x)$. Let \hat{x} and \hat{y} be the paths of Corollary 1. Since x and y are complements, $w(x) + w(y) = n$, so at least one of $w(x)$ or $w(y)$ (say $w(x)$) is less than or equal to $|n/2|$. Then the length of \hat{x} is bounded by $n + |n/2| = |3/2 n|$.

We now consider paths of the form \tilde{x} .

Lemma 3 Let x , s , and d be as in Lemma 2. If $\tilde{x}(s) = d$, then

$$d_i = \begin{cases} s_{n-1} \oplus x_0 \oplus x_1 & \text{if } i = 0 \\ s_{i-1} \oplus x_i \oplus x_{i+1} \oplus 1 & \text{if } 1 \leq i < n. \end{cases}$$

Proof Similar to Lemma 2.

Lemma 3 is illustrated for the case $n = 4$ in Figure 2.

Corollary 3 If $\zeta(d_i \oplus s_{i-1} \oplus 1) = 1$, then there exist strings $x, y \in 2^n$ such that $\tilde{x}(s) = \tilde{y}(s) = d$ and $x = \bar{y}$.

Proof For fixed s and d , Lemma 3 produces n equations in the n unknowns x_i :

$$\begin{aligned} x_1 &= d_0 \oplus s_{n-1} \oplus x_0 \\ x_{i+1} &= d_i \oplus s_{i-1} \oplus 1 \oplus x_i, \quad 0 \leq i < n. \end{aligned} \tag{3}$$

Equations (3), like equations (2), are not independent. Solving in terms of x_0 , we get

$$x_i = 1 \oplus x_0 \oplus \sum_{j < i} (d_j \oplus s_{j-1} \oplus 1), \quad 1 \leq i < n$$

$$\text{and } x_0 = 1 \oplus x_0 \oplus \sum_{i < n} (d_i \oplus s_{i-1} \oplus 1).$$

In this case, the equations have a solution iff $\sum (d_i \oplus s_{i-1} \oplus 1) = 1$. If so, the assignments $x_0 = 0$ and $x_0 = 1$ yield two complementary solutions to (3).

Corollary 4 If $\sum (d_i \oplus s_{i-1} \oplus 1) = 1$ then there is a path from s to d of length less than or equal to $\lfloor (3n-2)/2 \rfloor$.

Proof The length of the path \tilde{x} is $n-1+w(x)$. Let \tilde{x} and \tilde{y} be the paths of corollary 3. Since x and y are complements, at least one, say x , has weight $w(x) \leq \lfloor n/2 \rfloor$. Then the length of path \tilde{x} is bounded by $n-1+\lfloor n/2 \rfloor = \lfloor (3n-2)/2 \rfloor$.

Theorem The diameter of the Moebius graph of order n is less than or equal to $\lfloor 3/2 n \rfloor$.

Proof Let s and d be arbitrary vertices. There is a path from s to d of length $\leq \lfloor 3/2 n \rfloor$ if $\sum(d_i \oplus s_i \oplus 1) = 0$ (by Corollary 2), or if $\sum(d_i \oplus s_{i-1} \oplus 1) = 1$ by (Corollary 4). Since the sum in each case ranges over all values of i and subscripts are calculated modulo n , the sums are in fact identical and hence must be either both 0 or both 1.

The observation that $\sum(d_i \oplus s_i \oplus 1)$ is the parity of the number of positions in which d_i and s_i agree leads to the following simple algorithm for calculating paths:

Algorithm

Input: A pair of vertices s and d .

Output: A path from s to d of length $\leq \lfloor 3/2 n \rfloor$.

Method: If s and d agree in an even number of positions, define $x_0 = 0$ and for $i = 0, \dots, n-2$

$$x_{i+1} = s_i \oplus d_i \oplus 1 \oplus x_i.$$

If $w(x) \leq \lfloor n/2 \rfloor$, the desired path is \tilde{x} , otherwise it is \tilde{y} , where $y = \overline{x}$.

If s and d agree in an odd number of positions, let $x_0 = 0$, $x_1 = d_0 \oplus s_{n-1} \oplus x_0$, and for $i = 1, \dots, n-2$, let

$$x_{i+1} = d_i \oplus s_{i-1} \oplus 1 \oplus x_i.$$

If $w(x) \leq \lfloor n/2 \rfloor$, the desired path is \tilde{x} , otherwise it is \tilde{y} , where $y = \overline{x}$.

CONCLUSIONS AND OPEN PROBLEMS

We have described a family of undirected trivalent graphs having, for each n , $N = 2^n$ vertices and diameter bounded by $\lfloor 3/2 n \rfloor$. We are unable to determine the exact diameter, but for values of n up to 11 the actual diameter is $\lceil 3/2 n \rceil - 2$. More importantly, although the algorithm of the previous section calculates an adequate path between any pair of vertices, we know of no simple algorithm for calculating an optimal path. A related open question is the mean distance between pairs of vertices.

As we noted in the introduction, the Moebius graphs have an exponentially better density than any other infinite family of cubic graphs of which we are aware. However, they are still far from optimal. The only known upper bound on the number of vertices that can be packed within diameter k is the Moore bound derived as follows: Choose any vertex. There are at most $3 \cdot 2^{k-1}$ vertices at distance k from it. Hence the total number of vertices within distance k of the chosen vertex is $1 + \sum_{i \leq k} 3 \cdot 2^{i-1} = 3 \cdot 2^k - 2$. It has been shown [10,11] that this bound is only attained by two trivalent graphs: the complete graph K_4 with 4 vertices and diameter 1 and the Peterson graph with 10 vertices and diameter 2. For diameter 3, Elspas [12] gives the densest possible cubic graph (with 20 vertices, two less than the Moore bound). However, we are unaware of any general results indicating how far the actual optimum falls short of the Moore bound. Using a variety of techniques, we have obtained specific

trivalent graphs of diameters 4 through 10 with 34, 56, 84, 129, 210, 325, and 525 vertices, respectively [13]; in each case the diameter is bounded by $1.1 \lg N$. A much more elaborate form of the Moebius construction gives a family of trivalent graphs with diameter $1.47 \lg N + O(1)$ [14].

Finally, this paper only considers trivalent graphs. We are currently studying graphs of higher degrees.

5. REFERENCES

- [1] W. A. Wulf and C. G. Bell, "C.mmp -- a multi-miniprocessor," Proc. AFIPS 1972 Fall Joint Computer Conference 41, Part II, pp. 765-777 (1972).
- [2] R. M. Metcalfe and D. R. Boggs, "Ethernet: Distributed packet switching for local computer networks," CACM 19, 7, pp. 395-403 (July 1976).
- [3] D. J. Farber, J. Feldman, F. R. Heinrich, M. D. Hopwood, K. C. Larson, D. C. Loomis, and L. A. Rowee, "The Distributed Computing System," Proc. 7th Annual IEEE Computer Society International Conference, pp. 31-34 (February 1973).
- [4] M. H. Solomon and R. A. Finkel, "The Roscoe distributed operating system," Proc. 7th Symposium on Operating Systems Principles, pp. 108-114 (December 1979).
- [5] F. Harary, Graph Theory, Addison-Wesley, Reading, Massachusetts (1969).
- [6] B. W. Arden and H. Lee, "Analysis of chordal ring network," Proc. Workshop on Interconnection Networks for Parallel and Distributed Processing, pp. 93-100 (April 1980).
- [7] F. P. Preparata and J. Vuillemin, "The cube-connected-cycles: A versatile network for parallel computation," CACM 24, 5, pp. 300-309 (May 1981).

- [8] B. E. Arden and H. Lee, "A multi-tree-structured network," Proceedings of Compcon, Computer Communications Networks, pp. 201-210 (Fall 1978).
- [9] H. S. Stone, "Parallel processing with the perfect shuffle," IEEE Transactions on Computers C-20, 2, pp. 153-161 (February 1971).
- [10] E. Bannai and T. Ito, "On finite Moore graphs," Journal of the Faculty of Science of the University of Tokyo, Section IA 20, 2, pp. 191-208 (August 1973).
- [11] R. M. Damerell, "On Moore graphs," Proc. Cambridge Philosophical Society 74, 2, pp. 227-236 (September 1973).
- [12] B. Elspas, "Topological constraints on interconnection-limited logic," Switching Circuit Theory and Logical Design, pp. 133-137 (November 1964).
- [13] W. E. Leland, R. A. Finkel, Li Qiao, M. H. Solomon, and L. Uhr, "High density graphs for processor interconnection," Information Processing Letters 12, 3, pp. 117-120 (13 June 1981).
- [14] W. E. Leland, Dense interconnection graphs with small degree, Submitted for publication (August, 1981).

s_0	s_1	s_2	s_3
f			
s_1	s_2	s_3	$s_0 \oplus 1$
x_0			
s_1	s_2	$s_3 \oplus x_0$	$s_0 \oplus x_0 \oplus 1$
f			
s_2	$s_3 \oplus x_0$	$s_0 \oplus x_0 \oplus 1$	$s_1 \oplus 1$
x_1			
s_2	$s_3 \oplus x_0$	$s_0 \oplus x_0 \oplus x_1 \oplus 1$	$s_1 \oplus x_1 \oplus 1$
f			
$s_3 \oplus x_0$	$s_0 \oplus x_0 \oplus x_1 \oplus 1$	$s_1 \oplus x_1 \oplus 1$	$s_2 \oplus 1$
x_2			
$s_3 \oplus x_0$	$s_0 \oplus x_0 \oplus x_1 \oplus 1$	$s_1 \oplus x_1 \oplus x_2 \oplus 1$	$s_2 \oplus x_2 \oplus 1$
f			
$s_0 \oplus x_0 \oplus x_1 \oplus 1$	$s_1 \oplus x_1 \oplus x_2 \oplus 1$	$s_2 \oplus x_2 \oplus 1$	$s_3 \oplus x_0 \oplus 1$
x_3			
$s_0 \oplus x_0 \oplus x_1 \oplus 1$	$s_1 \oplus x_1 \oplus x_2 \oplus 1$	$s_2 \oplus x_2 \oplus x_3 \oplus 1$	$s_3 \oplus x_3 \oplus x_0 \oplus 1$

Fig. 1. Calculation of $\hat{x}(s)$

	s_0	s_1	s_2	s_3
x_0	s_0			
f		s_1	$s_2 \oplus x_0$	$s_3 \oplus x_0$
x_1	s_1	$s_2 \oplus x_0$	$s_3 \oplus x_0$	$s_0 \oplus 1$
f		$s_2 \oplus x_0$	$s_3 \oplus x_0 \oplus x_1$	$s_0 \oplus x_1 \oplus 1$
x_2	$s_2 \oplus x_0$	$s_3 \oplus x_0 \oplus x_1$	$s_0 \oplus x_1 \oplus 1$	$s_1 \oplus 1$
f		$s_2 \oplus x_0$	$s_3 \oplus x_0 \oplus x_1$	$s_1 \oplus x_2 \oplus 1$
x_3	$s_3 \oplus x_0 \oplus x_1$	$s_0 \oplus x_1 \oplus x_2 \oplus 1$	$s_1 \oplus x_2 \oplus 1$	$s_2 \oplus x_0 \oplus 1$
	$s_3 \oplus x_0 \oplus x_1$	$s_0 \oplus x_1 \oplus x_2 \oplus 1$	$s_1 \oplus x_2 \oplus x_3 \oplus 1$	$s_2 \oplus x_3 \oplus x_0 \oplus 1$

Fig. 2. Calculation of $\tilde{x}(s)$

END

FILMED

10-85

DTIC